Detail Matters: High-Frequency Content for Realistic Synthetic MRI Generation

Filip Rusak[1,2(✉)], Rodrigo Santa Cruz[2], Elliot Smith[3], Jurgen Fripp[2], Clinton Fookes[1], Pierrick Bourgeat[2], and Andrew Bradley[1]

[1] Queensland University of Technology, Brisbane, QLD, Australia
rus173@csiro.au
[2] CSIRO, Herston, QLD, Australia
[3] Maxwell Plus, Brisbane, QLD, Australia

Abstract. Deep Learning (DL)-based segmentation methods have been quite successful in various medical imaging applications. The main bottleneck of these methods is the scarcity of quality-labelled samples needed for their training. The lack of labelled training data is often addressed by augmentation methods, which aim to synthesise realistic samples with corresponding labels. While the synthesis of realistic samples remains a challenging task, little is known about the impact of fine detail in synthetic data on the performance of DL-based segmentation models. In this work, we investigate whether, and to what extent, the high-frequency (HF) detail in synthetic brain MR images (MRIs) impacts the performance of DL-based segmentation methods. To assess the impact of HF detail, we generate two synthetic datasets, with and without HF detail and train corresponding segmentation models to evaluate the impact on their performance. The results obtained demonstrate that the presence of HF detail in synthetic brain MRIs, used during training, significantly improve the Dice score up to 1.73% for Gray Matter (GM), 1.34% for White Matter (WM) and 4.41% for Cerebrospinal Fluid (CSF); and therefore justify the need for synthesising realistic-looking MRIs.

Keywords: Data augmentation · Brain MRI · Generative adversarial network · Realistic brain MRI synthesis

1 Introduction

In recent years, Convolutional Neural Networks (CNNs) have been successfully applied on segmentation tasks in the field of medical imaging [5,6,9]. The quality of segmentation results largely depends on the training method, data and, in some cases, corresponding labels required for the segmentation model training [19]. Supervised segmentation methods perform better in comparison to the other

This work was funded in part through an Australian Department of Industry, Energy and Resources CRC-P project between CSIRO, Maxwell Plus and I-Med Radiology Network.

D. Svoboda et al. (Eds.): SASHIMI 2021, LNCS 12965, pp. 3–13, 2021.
https://doi.org/10.1007/978-3-030-87592-3_1

types of learning, but rely on large amounts of quality labelled data which is often not available [16,30]. The reason for a large train set requirement is to provide CNNs with diverse samples, so they can generalise and perform segmentation tasks on unseen samples [10].

Data augmentation is commonly used to bridge the gap between the available and desired quantity of samples, which is reflected as an improvement in segmentation results [7,30]. Data augmentation methods can be broadly grouped as traditional, mixing or synthetic [22]. Traditional augmentation methods such as scaling, rotation, flipping and translation provide limited variation improvements. Mixup [29] in recent years, due to its simplicity, gained popularity in the context of data augmentation for classification, but it has limited applicability on segmentation tasks with sometimes mixed results [4]. Synthetic augmentation methods, for uni or cross-domain scenarios, are commonly realised in some form of Generative Adversarial Networks (GANs) [8]. GANs are capable of producing more diverse samples than traditional augmentation methods, resulting in the improvement of model generalisation ability [28]. Further, conditional GANs [17] and their variations [11,15,31] provide more control over data synthesis and, in the context of brain MR imaging, allow the generation of brain MRIs with particular anatomy, pathology or modality.

Recent work [32] addressed the problem of unrealistic-looking synthetic brain MRI generation (commonly over-smoothed without fine detail) and explained it as a manifestation of distribution mismatch between real and synthetic scans. Intuitively, such an argument motivates the generation of realistic brain MRI samples which remains a challenging task. Challenges in the generation of realistic brain MRI scans come from the versatile nature of MRI representations, 3D shape and the scan sizes [13,23]. The major technical hurdle related to realistic synthetic brain MRIs generation is the limited GPU memory [21,23]. Such an obstacle was previously tackled by working with slices or patches [23], employing series of GANs [21,23] or splitting the generation of shape and texture/appearance [14,30] into separate tasks. The aforementioned methods obtained good results, but often use complex architectures, require training of multiple models or lack control over generated samples. Further, while state-of-the-art methods have focused on the synthesis of realistic medical images [14,21,23,32], little is known about the effect of synthetic image realism on the MRI brain segmentation models.

In this paper, we evaluate the effect of high-frequency (HF) detail, as a single aspect of synthetic brain MRI realism, on the performance of brain segmentation methods. First, we propose a single GAN-based method to generate 3D synthetic MRIs with HF detail. Second, we generate two synthetic brain MRI datasets, with and without HF detail. Then, we evaluate the impact of HF detail on segmentation performance, by training segmentation models on three datasets: i) real MRIs, ii) real + synthetic MRIs with and iii) without HF detail. The obtained results suggest that synthetic can replace real MRIs for the purpose of model training while achieving comparable results. Finally, we conclude that

HF detail matters for the training of brain segmentation models and therefore justify the efforts needed to generate synthetic brain MRIs with HF detail.

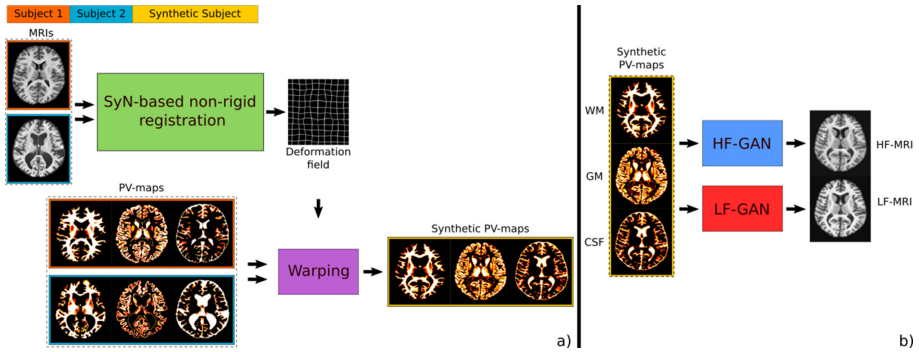

Fig. 1. Non-rigid registration-based MRI and PV-map augmentation (a), LF and HF synthetic brain MRI generation (b).

2 Methods

Data and Pre-processing. 1000 T1-w brain MRI scans from ADNI[1] [12,27] were used to train, validate and test, both, brain MRI synthesis and segmentation models. For the purpose of training, validation and testing of the GAN model, the 1000 MRI scans were split randomly in the ratio 60:20:20, respectively, without any overlaps between the sets, while retaining the equal distribution of gender and diagnosis. The identical validation and test sets were used for GAN and segmentation models. In the case of segmentation models, the train set was created from 250/600 randomly selected scans included in the GAN models train set. All scans were pre-processed using bias field correction in the brain region of interest (ROI) [24], rigid registration to the MNI-space ($181 \times 217 \times 181$ voxels) and z-score intensity normalisation with the mean value computed from brain ROI. The corresponding labels were derived by a segmentation method implemented based on the expectation-maximisation algorithm [25]. Further, Partial Volume (PV)-maps were estimated from derived labels following the PV-estimation algorithm described in [1].

[1] Data used in the preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). The ADNI was launched in 2003 as a public-private partnership, led by Principal Investigator Michael W. Weiner,MD. The primary goal of ADNI has been to test whether serial magnetic resonance imaging (MRI), positron emission tomography (PET), other biological markers, and clinical and neuropsychological assessment can be combined to measure the progression of mild cognitive impairment (MCI) and early Alzheimer's disease (AD). For up-to-date information, see www.adni-info.org.

Brain MRI Synthesis. In this section, we focus on the generation of synthetic MRI scans from three tissue classes represented with PV-maps: WM, GM and CSF. We define a PV-map (M_{pv}) as a volume $M_{pv} \in [0,1]^{w \times h \times d}$ that corresponds to a particular MRI scan and represents segmentation of a certain tissue class with sub-voxel precision. The sub-voxel precision is achieved by assigning a value within [0,1] interval to each voxel. The voxel value stands for the proportion of a particular tissue type in a voxel. To generate novel synthetic brain MRI scans that correspond to a subject in-between two real subjects, we use non-rigid registration between two selected baseline subject MRIs. The symmetric image normalization method (SyN) [2], implemented in the Advanced Normalization Tools [3] package, is used to register two baseline subject MRIs. The half-way deformation field (middle of the normalisation domain) is then used to warp the corresponding PV-maps. As a result, we obtained PV-maps of a synthetic subject as shown in Fig. 1(a). Once synthetic subject PV-maps are created, we used them to generate synthetic MRI scans with and without HF detail as shown in Fig. 1(b).

A conditional GAN-based solution was selected for the generation of more and less realistic synthetic MRIs from PV-maps of three tissues (WM, GM and CSF). We trained the same GAN in two stages, aiming to obtain two models able to generate synthetic brain MRIs of both levels of realism (Fig. 3). For ease of reference, a GAN model capable of brain MRI synthesis without HF detail, from now on will be referred to as LF-GAN. In contrast, a GAN model capable of brain MRI synthesis with HF detail will be referred to as HF-GAN. The model architecture is based on pix2pix framework [11] and is implemented according to the architecture provided in [20] with further adjustments in the second training stage (HF-GAN model). The detailed architecture components used in our implementation is shown in Fig. 2. Both models employed the same U-Net [18] based generator.

The main difference between the first (LF-GAN) and the second (HF-GAN) training stage is in the discriminator being used. In the first training stage, we train a model according to the architecture outlined in [20]. A PatchGAN discriminator, used in [20], is known to be limited in recovering detail on different scales [26], which allows us to obtain synthetic brain MRIs without HF detail. We trained the model, with the stopping criteria defined as generator loss plateaus for at least 10 epochs with fluctuations no more than 0.01. In the second training stage, we resumed the training with the pre-trained generator and the ResNet discriminator, typically used in super-resolution (SR) applications [15]. Training a GAN, with a ResNet discriminator, without pre-training the generator leads to poor results. In SR applications, the input to SR-GANs are low-resolution (LR) images and the outputs are synthetic images with estimated high-resolution detail. In our case, the resolution of the input and output MRIs remains the same but the input MRIs, generated by a pre-trained generator, visually resemble LR images which allow further recovery of HF-detail during the second training stage. The same training criteria was used for both training stages.

To formalise the training of both GAN models, let data of a certain distribution d_x be denoted with x, a generator with G, its output $G(c_{1-3}, z)$ and a discriminator with D. Moreover, we denote three condition variables with c_{1-3}(M_{pv} for three tissue-types) and a noise variable with z. The objective function is defined as follows,

$$\min_G \max_D \mathbb{E}_{c_{1-3},x}\left[log\left(D\left(c_{1-3},x\right)\right)\right] + \mathbb{E}_{c_{1-3},z}\left[log\left(1-D\left(c_{1-3},G\left(c_{1-3},z\right)\right)\right)\right] + \mathbb{E}_{c_{1-3},x,z}\left[\left\|x-G\left(c_{1-3},z\right)\right\|_1\right]. \tag{1}$$

To evaluate the benefit of synthetic HF-MRIs in the context of brain segmentation model training, we trained 15 brain segmentation models, split into three categories according to the train set composition (Real, Real + LF and Real + HF synthetic MRIs). The data split across the models is outlined in Table 2. The primary aim of this experiment is to evaluate the benefits of using HF synthetic brain MRIs for the training of brain segmentation models rather than achieve a Dice score higher than state-of-the-art brain segmentation methods.

GAN Training. Both models were initialised by Adam optimiser and trained, until satisfying the stopping criteria, with the batch size set to one due to the memory limit and the image size. The initial learning rate of 0.0002 was for training both models. The learning rate decay of 2×10^{-6} was introduced in HF-GAN model after 20^{th} epoch.

Segmentation. The segmentation model used for the evaluation of synthesised MRIs is based on vox2vox[2] [5] which was adapted to facilitate the segmentation of four classes: WM, GM, CSF and background. The input to the model is a skull-stripped brain MRI scan and the model generates a binary segmentation map for each of four classes. To fit in the architecture of the segmentation model, the scans and the corresponding binary labels were zero-padded to the size of $256 \times 256 \times 256$ voxels. The binary labels were created from PV-maps by assigning a voxel to the class with the highest PV-value at the same location across three PV-maps (WM, GM and CSF).

Table 1. Image quality assessment by full-reference IQM (left) and segmentation/PV-estimation error of three tissue classes for LF and HF relative to real MRIs (right). The * denotes statistically significant results (t-test after Holm-Bonferroni correction).

Dataset	MS-SSIM ↑	NRMSE ↓	PSNR ↑	NRMSE ↓		
				GM	WM	CSF
LF	0.9834 ± 0.0085	0.0121 ± 0.0423	34.7167 ± 1.9913	0.04343 ± 0.1893	0.0321 ± 0.1465	0.0554 ± 0.1627
HF	**0.9870 ± 0.0071***	**0.0118 ± 0.0437**	**36.5306 ± 1.8113***	**0.0428 ± 0.2806**	**0.0301 ± 0.2852**	**0.0546 ± 0.25092**

[2] The pytorch implementation of vox2vox taken from https://github.com/enochkan/vox2vox.

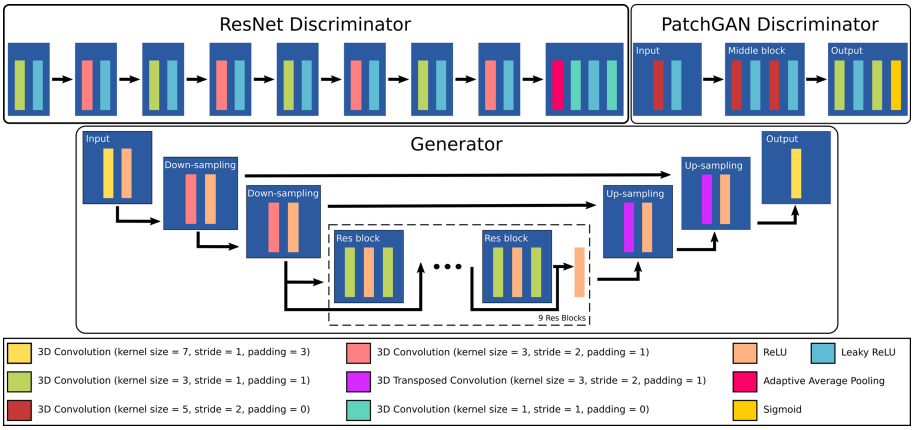

Fig. 2. Architectural building-blocks used for both LF and HF-GANs, and complementary implementation detail.

3 Experiments and Results

Quantitative Evaluation of Synthetic MRIs. To evaluate the synthetic scan quality at both levels of realism (LF vs. HF scans), we reconstructed LF and HF synthetic MRIs from the PV-maps and measured their image quality in relation to the corresponding real MRIs. The PV-maps used for the MRI synthesis were derived from the real MRIs included in the test set (200 samples). Once the synthetic MRIs were generated, the image quality was then measured, between the real and both types of synthetic MRIs, using the following full-reference Image Quality Metrics (IQM): Multiscale Structural Similarity Index Measure (MS-SSIM), Normalised Root Mean Square Error (NRMSE) and Peak Signal-to-Noise Ratio (PSNR). The results, presented in the Table 1 (left), indicate that HF-MRIs have higher similarity with real MRIs compared to the LF-MRI scans. Further, we derived PV-maps from real and synthesised MRIs, and for each tissue class (GM, WM and CSF), we computed the NRMSE between the PV-maps derived from real and both types of synthetic MRIs. The rationale behind that metric suggests that the difference in PV-maps derived from more similar MRIs should be smaller than in the case of the less similar MRIs. The tissue class-wise results, presented in Table 1 (right), indicate a higher similarity between real and HF than between real and LF-MRIs. According to the results presented in the Table 1, the computed IQMs numerically support the higher visual similarity between real and HF synthetic MRIs, compared to the LF synthetic MRIs, as illustrated in Fig. 3.

The Impact of HF Detail on Brain Segmentation. After training the aforementioned segmentation models, we segmented three tissue-classes (GM, WM and CSF) from each of 200 test real MRI samples with each of the trained segmentation models. The model performance was measured using Dice score. The

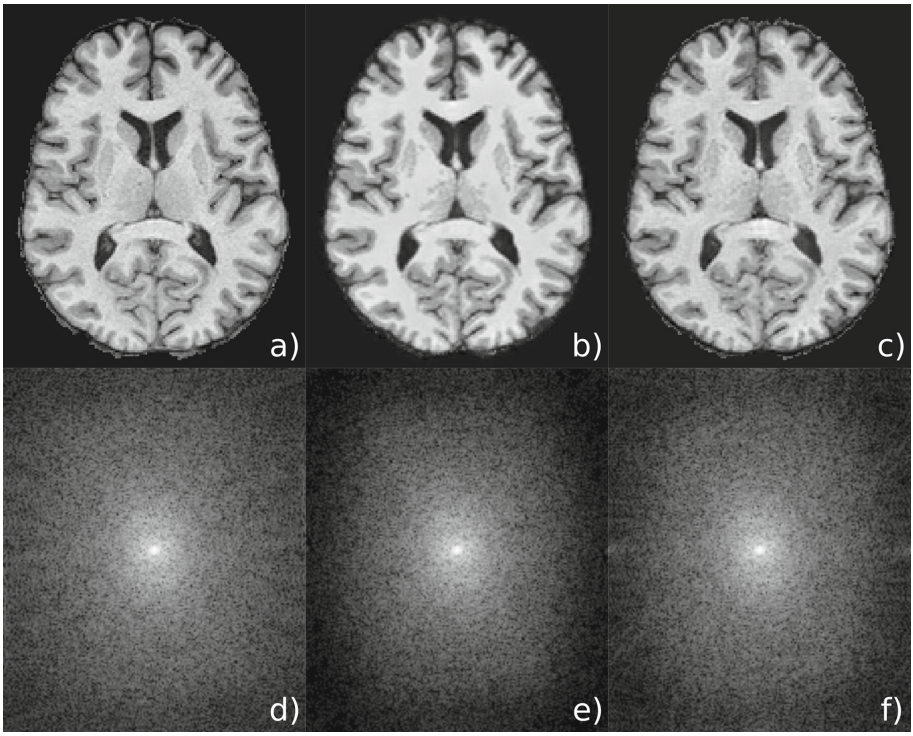

Fig. 3. Real brain MRI (a), synthetic scan generated by LF-GAN (b) and synthetic scan generated by HF-GAN. The respective representation in Fourier domain shown on (d), (e) and (f).

measurement results are presented in the Table 2 and Fig. 4. The upper section of Table 2 shows the segmentation results obtained from five models trained on the real MRIs and corresponding segmentation of the three tissue classes. The bottom section shows the results obtained from the eight remaining segmentation models. According to the results presented in Table 2 and Fig. 4, both LF and HF-MRIs, improve the segmentation model performance for all three tissue classes. It is also evident that the proposed augmentation method can achieve comparable performance to models trained on real MRIs only, with only 20% of real samples. When comparing the overall segmentation performance of models trained on a mixed dataset (real + synthetic MRIs), the models trained on a HF dataset (real + HF synthetic MRIs) perform better than models trained on a LF dataset (real + LF synthetic MRIs). In the case of GM segmentation, the models trained on 50 real and 100 as well as 150 HF synthetic MRIs shows statistically significant improvement in comparison to models trained on the same number of LF synthetic MRI samples. In the case of WM segmentation, models trained on a mixed dataset of 50 real and 50, 100 and 150 HF synthetic MRIs as well as 250 real and 200 HF synthetic show statistically significant improvement in

comparison to models trained on the same number of LF synthetic MRI samples. The models trained on 50 real MRIs and 100 and 150 synthetic HF MRIs achieve almost the same performance, in WM segmentation, as the models trained on real MRIs of the same data sample size. In the case of CSF segmentation, all models trained on a HF datasets, except the last model (250 real & 200 synthetic MRIs), show statistically significant improvement comparing to models trained on LF datasets.

The largest Dice score improvement of models trained on HF, compared to, LF dataset, regarding GM segmentation is 1.73%, WM segmentation 1.34%, and CSF segmentation is 4.41%.

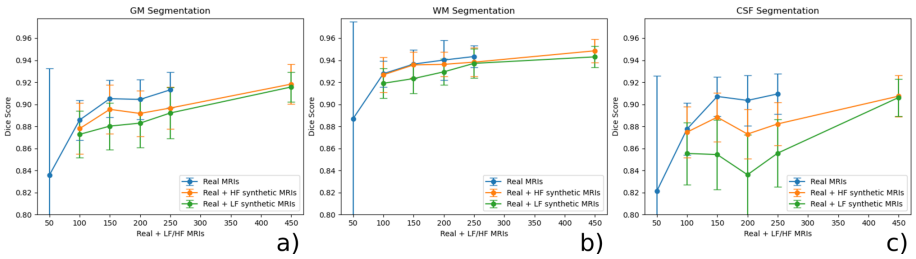

Fig. 4. Performance comparison of segmentation models, represented as a mean Dice score and standard deviation ($\pm\sigma$), trained on real and LF/HF synthetic MRIs, for each of three tissue classes: GM (a), WM (b) and CSF (c).

Table 2. Dice scores of 15 segmentation models trained on datasets with different data splits. The * denotes statistically significant results (t-test after Holm-Bonferroni correction).

Data split (real : synthetic)	GM		WM		CSF	
50:0	0.8356 ± 0.0882		0.8869 ± 0.0967		0.8211 ± 0.1046	
100:0	0.8857 ± 0.018		0.9277 ± 0.0118		0.8775 ± 0.0236	
150:0	0.9051 ± 0.0169		0.9364 ± 0.0132		0.9071 ± 0.0178	
200:0	0.9044 ± 0.0182		0.9401 ± 0.0182		0.9035 ± 0.0228	
250:0	0.9131 ± 0.0161		0.9433 ± 0.01		0.9094 ± 0.0184	
	LF	HF	LF	HF	LF	HF
50:50	0.8727 ± 0.0213	**0.8779 ± 0.0231**	0.919 ± 0.0135	**0.9267 ± 0.016***	0.8554 ± 0.0282	**0.8747 ± 0.0231***
50:100	0.8802 ± 0.0212	**0.8954 ± 0.022***	0.9233 ± 0.0133	**0.9357 ± 0.0116***	0.8544 ± 0.0316	**0.8883 ± 0.0221***
50:150	0.8829 ± 0.0223	**0.8917 ± 0.0208***	0.9294 ± 0.012	**0.9362 ± 0.0111***	0.8361 ± 0.0364	**0.873 ± 0.0224***
50:200	0.8921 ± 0.0229	**0.8966 ± 0.0189**	0.9371 ± 0.0131	**0.9382 ± 0.0129**	0.8557 ± 0.0305	**0.8822 ± 0.0194***
250:200	0.9157 ± 0.0135	**0.9183 ± 0.0180**	0.943 ± 0.0097	**0.9485 ± 0.0106***	0.9062 ± 0.0169	**0.9074 ± 0.0188**

4 Conclusion

In this paper, we proposed a simple, yet effective, GAN training method for the generation of realistic-looking synthetic MRIs and evaluated the impact of

HF detail on the performance of a DL-based brain segmentation method. The results suggest that, in the context of segmentation model training, real data can be replaced by synthetic MRIs while still achieving comparable results. We further showed that synthetic MRIs with HF detail can significantly improve segmentation results on a downstream task, indicating that the realism of synthetic samples matters. For future work, we plan a comprehensive evaluation on the impact of HF detail on other state-of-the-art DL-based brain segmentation methods.

References

1. Acosta, O., et al.: Automated voxel-based 3D cortical thickness measurement in a combined Lagrangian-Eulerian PDE approach using partial volume maps. Med. Image Anal. **13**(5), 730–743 (2009)
2. Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C.: Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. Med. Image Anal. **12**(1), 26–41 (2008)
3. Avants, B.B., Tustison, N., Song, G.: Advanced normalization tools (ANTS). Insight J **2**(365), 1–35 (2009)
4. Chen, C., et al.: Realistic adversarial data augmentation for MR image segmentation. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12261, pp. 667–677. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59710-8_65
5. Cirillo, M.D., Abramian, D., Eklund, A.: Vox2vox: 3D-GAN for brain tumour segmentation. arXiv preprint arXiv:2003.13653 (2020)
6. Coupé, P., et al.: AssemblyNet: a large ensemble of CNNs for 3D whole brain MRI segmentation. Neuroimage **219**, 117026 (2020)
7. Eaton-Rosen, Z., Bragman, F., Ourselin, S., Cardoso, M.J.: Improving data augmentation for medical image segmentation (2018)
8. Goodfellow, I., et al.: Generative adversarial nets. In: Advances in neural information processing systems, pp. 2672–2680 (2014)
9. Henschel, L., Conjeti, S., Estrada, S., Diers, K., Fischl, B., Reuter, M.: FastSurfer-a fast and accurate deep learning based neuroimaging pipeline. NeuroImage 117012 (2020)
10. Huang, S.G., Chung, M.K., Qiu, A., Initiative, A.D.N.: Fast mesh data augmentation via chebyshev polynomial of spectral filtering. arXiv preprint arXiv:2010.02811 (2020)
11. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1125–1134 (2017)
12. Jack Jr., C.R., et al.: The Alzheimer's disease neuroimaging initiative (ADNI): MRI methods. J. Magn. Reson. Imaging: Off. J. Int. Soc. Magn. Reson. Med. **27**(4), 685–691 (2008)
13. Jog, A., Carass, A., Roy, S., Pham, D.L., Prince, J.L.: MR image synthesis by contrast learning on neighborhood ensembles. Med. Image Anal. **24**(1), 63–76 (2015)
14. Keong, C.C., Wei, H.E.T.: Synthesis of 3D MRI brain images with shape and texture generative adversarial deep neural networks. IEEE Access **9**, 64747–64760 (2021)

15. Ledig, C., et al.: Photo-realistic single image super-resolution using a generative adversarial network. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4681–4690 (2017)
16. Lee, J., Kim, E., Lee, S., Lee, J., Yoon, S.: FickleNet: weakly and semi-supervised semantic image segmentation using stochastic inference. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 5267–5276 (2019)
17. Mirza, M., Osindero, S.: Conditional generative adversarial nets. arXiv preprint arXiv:1411.1784 (2014)
18. Ronneberger, O., Fischer, P., Brox, T.: U-net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
19. Roy, A.G., Conjeti, S., Navab, N., Wachinger, C., Initiative, A.D.N., et al.: QuickNAT: a fully convolutional network for quick and accurate segmentation of neuroanatomy. Neuroimage **186**, 713–727 (2019)
20. Rusak, F., et al.: 3D brain MRI GAN-based synthesis conditioned on partial volume maps. In: Burgos, N., Svoboda, D., Wolterink, J.M., Zhao, C. (eds.) SASHIMI 2020. LNCS, vol. 12417, pp. 11–20. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59520-3_2
21. Sun, L., Chen, J., Xu, Y., Gong, M., Yu, K., Batmanghelich, K.: Hierarchical amortized training for memory-efficient high resolution 3D GAN. arXiv preprint arXiv:2008.01910 (2020)
22. Tajbakhsh, N., Jeyaseelan, L., Li, Q., Chiang, J.N., Wu, Z., Ding, X.: Embracing imperfect datasets: a review of deep learning solutions for medical image segmentation. Med. Image Anal. **63**, 101693 (2020)
23. Uzunova, H., Ehrhardt, J., Jacob, F., Frydrychowicz, A., Handels, H.: Multi-scale GANs for memory-efficient generation of high resolution medical images. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11769, pp. 112–120. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32226-7_13
24. Van Leemput, K., Maes, F., Vandermeulen, D., Suetens, P.: Automated model-based bias field correction of MR images of the brain. IEEE Trans. Med. Imaging **18**(10), 885–896 (1999)
25. Van Leemput, K., Maes, F., Vandermeulen, D., Suetens, P.: Automated model-based tissue classification of MR images of the brain. IEEE Trans. Med. Imaging **18**(10), 897–908 (1999)
26. Wang, J., Chen, Y., Wu, Y., Shi, J., Gee, J.: Enhanced generative adversarial network for 3D brain MRI super-resolution. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pp. 3627–3636 (2020)
27. Weiner, M.W., et al.: The Alzheimer's disease neuroimaging initiative 3: continued innovation for clinical trial improvement. Alzheimer's Dementia **13**(5), 561–571 (2017)
28. Yi, X., Walia, E., Babyn, P.: Generative adversarial network in medical imaging: a review. Med. Image Anal. **58**, 101552 (2019)
29. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: beyond empirical risk minimization. arXiv preprint arXiv:1710.09412 (2017)
30. Zhao, A., Balakrishnan, G., Durand, F., Guttag, J.V., Dalca, A.V.: Data augmentation using learned transformations for one-shot medical image segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 8543–8553 (2019)

31. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 2223–2232 (2017)
32. Zuo, L., et al.: Synthesizing realistic brain MR images with noise control. In: Burgos, N., Svoboda, D., Wolterink, J.M., Zhao, C. (eds.) SASHIMI 2020. LNCS, vol. 12417, pp. 21–31. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59520-3_3